arXiv:1510.00732v1 [math.LO] 2 Oct 2015

TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 00, Number 0, Pages 000–000
S 0002-9947(XX)0000-0

ZERO SETS OF UNIVARIATE POLYNOMIALS

ROBERT S. LUBARSKY AND FRED RICHMAN

Abstract. Let L be the zero set of a nonconstant monic polynomial with complex coefficients. In the context of constructive mathematics without countable choice, it may not be possible to construct an element of L. In this paper we introduce a notion of distance from a point to a subset, more general than the usual one, that allows us to measure distances to subsets like L. To verify the correctness of this notion, we show that the zero set of a polynomial cannot be empty—a weak fundamental theorem of algebra. We also show that the zero sets of two polynomials are a positive distance from each other if and only if the polynomials are comaximal. Finally, the zero set of a polynomial is used to construct a separable Riesz space, in which every element is normable, that has no Riesz homomorphism into the real numbers.

1. Quasidistance

Let T be a set of real numbers. A **lower bound** for T is a real number b such that $b \le t$ for all t in T. By the **infimum** of T we mean, as usual, a lower bound $\inf T$ for T such that if $\inf T < r$, then there exists $t \in T$ such that $t < r$. The number $\sup T$ is defined dually. By the **greatest lower bound** of T we mean the largest element, $\operatorname{glb} T$, of the set B of lower bounds of T. As the distinction between $\inf T$ and $\operatorname{glb} T$ in constructive mathematics is central to this section, we pause to illustrate this difference.

Example 1.1. Let P be an arbitrary proposition and let

$$L = \{0 : P \text{ or } \neg P\} \cup \{1\}\,.$$

Then the set B of lower bounds of L is $(-\infty, 0]$. Thus $\operatorname{glb} L = 0$, but if $\inf L$ exists, then P or $\neg P$.

Proof. Elements of $(-\infty, 0]$ are certainly lower bounds of L. Conversely, if $b > 0$, then b cannot be a lower bound of L, because if 0 were not in L, then P would have to be false, so $\neg P$ would be true, so 0 would be in L. So elements of B cannot be (strictly) positive, whence they are in $(-\infty, 0]$.

Finally, if $m = \inf L$, then $m \le 0$ so $1 > m$ whence there exists $r < 1$ in L. From this P or $\neg P$ follows. □

So the notions of infimum and greatest lower bound are the same only in the presence of the law of excluded middle. The relations that hold between the two notions are stated in the following theorem.

2000 *Mathematics Subject Classification.* Primary 03F65, 13A99.

©XXXX American Mathematical Society

Theorem 1.2. *Let T be a set of real numbers that is bounded below, and B the set of lower bounds for T. If $\sup B$ exists, then $\operatorname{glb} T = \sup B$. If $\inf T$ exists, then $\operatorname{glb} T = \inf T$.*

Proof. Suppose $\sup B$ exists. We must show that it is in B, that is, that $t \geq \sup B$ for all t in T. But $t \geq b$ for each b in B, so $t \geq \sup B$. (If $t < \sup B$, then there would be $b \in B$ such that $t < b$.)

Suppose $\inf T$ exists. We must show that it is the maximum element of B. It is in B by definition. If $b \in B$ and $b > \inf T$, then there exists t in T such that $t < b$, a contradiction, so $b \leq \inf T$. □

The usual informal description of a located subset L of a metric space M is that we can compute the distance from any point in M to L. This is slightly misleading because the rules for computing the distance are not explicitly stated. It is not enough to be able to compute the distance from x to L, we must compute the infimum

$$\inf_{y \in L} d(x, y)$$

Thus there are some implied constructions of elements of L hidden in that statement about computing distances. That is not the case for the greatest lower bound

$$\operatorname*{glb}_{y \in L} d(x, y)$$

which is also a candidate for the distance from x to L.

To illustrate this distinction, let M be the real line and L the subset of Example 1.1. The subset L contains an element, is closed, and $\operatorname{glb}_{y \in L} d(-1, y) = 1 + \operatorname{glb} L = 1$. However, $\inf_{y \in L} d(-1, y) = 1 + \inf L$, but L does not have an infimum.

If L is a subset of a metric space M, and x is a point in M, then the **distance** from x to L is

$$d(x, L) = \inf_{y \in L} d(x, y)$$

We define the **quasidistance** from x to L to be

$$\delta(x, L) = \operatorname*{glb}_{y \in L} d(x, y)$$

We say that L is **located** if $d(x, L)$ exists for each x in M, and that L is **quasilocated** if $\delta(x, L)$ exists for each x in M. It follows from Theorem 1.2 that a located set L is quasilocated with $\delta(x, L) = d(x, L)$ for all x. Note that a quasilocated set cannot be empty.

Theorem 1.3. *If L is a quasilocated subset of a metric space M, then $\delta(x, L) \leq \delta(y, L) + d(x, y)$ for all x and y in M. It follows that $|\delta(x, L) - \delta(y, L)| \leq d(x, y)$ so $\delta(x, L)$ is a uniformly continuous function of x.*

Proof. It suffices to show that $\delta(x, L) - d(x, y) \leq d(y, t)$ for all $t \in L$. But $d(x, t) \leq d(x, y) + d(y, t)$, and $\delta(x, L) \leq d(x, t)$. □

If we think of $\delta(x, L)$ as a lowercut (a supremum of a bounded set of real numbers), see [4], then the inequality $\delta(x, L) \leq \delta(y, L) + d(x, y)$ holds for an arbitrary subset L of M.

The quasidistance can be thought of as dual to the distance. That's because we can express the distance from x to L as

$$d(x,L) = \inf\{d(x,t) : t \in L\} = \inf\{r : d(x,t) \leq r \text{ for some } t \text{ in } L\}$$

while the quasidistance is

$$\delta(x,L) = \sup\{r : d(x,t) \geq r \text{ for all } t \text{ in } L\}$$

Throughout, we will use $B_r(x)$ to denote the open ball of radius r centered at x.

Lemma 1.4. *Let L be a subset of a metric space M and $x \in M$. Then r is a lower bound for $\{d(x,y) : y \in L\}$ if and only if $B_r(x) \cap L = \emptyset$.*

Proof. Consider the statement P that there exists $y \in L$ such that $d(x,y) < r$. Then $\neg P$ says $d(x,y) \geq r$ for all $y \in L$, that is, that r is a lower bound for $\{d(x,y) : y \in L\}$. On the other hand, $\neg P$ says that $B_r(x) \cap L = \emptyset$. □

Theorem 1.5. *If $\delta(x,L)$ exists, then $\delta(x,L) \leq r$ if and only if $B_s(x) \cap L$ cannot be empty for any $s > r$.*

Proof. Suppose $\delta(x,L) = \text{glb}_{y \in L} d(x,y) \leq r$. If $B_s(x) \cap L$ were empty for $s > r$, then s would be a lower bound for $\{d(x,y) : y \in L\}$, a contradiction. Conversely, suppose $B_s(x) \cap L$ cannot be empty for any $s > r$. If $\delta(x,L) > r$, then $B_s(x) \cap L$ is empty for $s = (\delta(x,L) + r)/2$, a contradiction. □

We really aren't particularly interested in Example 1.1. What motivated these ideas was the study, in a choiceless environment, of the set L of roots of a monic polynomial p with complex coefficients. If the coefficients of p are complex numbers over the Cauchy reals, then we can write $p(X) = (X - r_1) \cdots (X - r_n)$, see [5], and L is the closure of the set $\{r_1, \ldots, r_n\}$, so L is easily seen to be located. However, if the coefficients are arbitrary complex numbers, then we need not be able to construct a root of p, see [3], even if p is of the form $X^2 - a$, so we can't show that L is located. However, L is always quasilocated, as we will see (Corollary 4.6).

In the presence of choice, at least in complete spaces M where balls are totally bounded (such as Euclidean spaces), quasilocated zero sets are located. To see that some hypothesis on M is necessary to get this conclusion, take L as in Example 1.1, and $M = \{-1\} \cup L$. Then L is the zero set of the function $x(x-1)$ on M.

Theorem 1.6. *Let M be a complete space in which balls are totally bounded. Assuming the axiom of dependent choices, quasilocated zero sets of uniformly continuous functions on M are located.*

Proof. Suppose L is a quasilocated subset of M and $x \in M$. We will show that if $\delta(x,L) < r$, and $\varepsilon > 0$, then there exists $t \in L$ such that $d(x,t) \leq r + \varepsilon$. The proof is reminiscent of the proof of the Baire category theorem. First we construct an element $y_1 \in B_r(x)$ such that $\delta(y_1,L) < \varepsilon/2$. This is possible because the function $\delta(z,L)$ is uniformly continuous, so has an infimum m as z ranges over $B_r(x)$. Either $m > 0$ or $m < \varepsilon/2$. But if $m > 0$, then $\delta(x,L) \geq r$, so $m < \varepsilon/2$.

Similarly we construct an element $y_2 \in B_{\varepsilon/2}(y_1)$ such that $\delta(y_2,L) < \varepsilon/4$. Then an element $y_3 \in B_{\varepsilon/4}(y_2)$ such that $\delta(y_3,L) < \varepsilon/8$. In general, we construct $y_{n+1} \in B_{\varepsilon/2^n}(y_n)$ such that $\delta(y_{n+1},L) < \varepsilon/2^{n+1}$. Dependent choice allows us to form a sequence $y_1, y_2, \ldots$ with those properties.

Because $y_{n+1} \in B_{\varepsilon/2^n}(y_n)$, the sequence of y's is a Cauchy sequence, hence converges to a point t in M. As $\delta(y_n, L)$ converges to zero, $\delta(t, L) = 0$. As L is a zero set of a uniformly continuous function, $t \in L$. Finally $d(x, y_n) < r + \varepsilon$ for all n, so $d(x, t) \leq r + \varepsilon$. □

With a little work, the proof of this theorem can be modified to apply to any locally compact space M.

In Example 1.1, the quasidistance from the point 0 to the closed set L is equal to zero but 0 need not be in L. That can't happen if L is the zero set of a continuous function. We write this as a lemma for later use.

Lemma 1.7. *Let f be a continuous function that vanishes on L. If $\delta(x, L) = 0$, then $f(x) = 0$.*

Proof. It's enough to show that $f(x)$ cannot be different from 0. Suppose $|f(x)| = \varepsilon > 0$. As f is continuous at x, there is $\theta > 0$ such that f is nonzero in $B_\theta(x)$. Lemma 1.4 then tells us that $\delta(x, L) \geq \theta$, a contradiction. □

2. Quasifinite sets

Even with choice, the zero set of a monic polynomial need not be finite or even finitely enumerable. For example, let a and b are two real numbers and consider the polynomial $p(X) = (X - a)(X - b)$. In addition to the roots a and b, the numbers $a \vee b$ and $a \wedge b$ are roots, and the ability to equate $a \vee b$ with either a or b is a form of LLPO. What we can say is that the zero set of $p(X)$ is the closure of the set $\{a, b\}$. So what we are thinking of here, in the presence of choice, are closures of finitely enumerable sets.

What are we thinking of in the absence of choice? The example to keep in mind is the zero set L of the polynomial $X^2 - a$ where a is a small complex number, possibly zero. Unless there is a sequence of Gaussian numbers converging to a, see [5], we cannot necessarily construct an element of L, see [3]. But L is quasilocated, as we shall see. First a little terminology.

An **ε-quasiapproximation** to a subset L of a metric space M is a sequence $x_1, \ldots, x_n$ in M such that $B_\varepsilon(x_i) \cap L$ cannot be empty for $i = 1, \ldots, n$, and $L \subset B_\varepsilon(x_1) \cup \cdots \cup B_\varepsilon(x_n)$. We say that L is **quasi-totally-bounded** if for each $\varepsilon > 0$ we can find an ε-quasiapproximation to L. Unlike the notion of totally bounded, this depends on the ambient space M because we cannot assume that the ε-quasiapproximation lies in L. The subset L in Example 1.1 is quasi-totally-bounded but we can't find an internal 1-quasiapproximation to it.

Theorem 2.1. *Each quasi-totally-bounded subset of a metric space M is quasilocated.*

Proof. Let L be a quasi-totally-bounded subset of M. For any $\varepsilon > 0$, let $x_1, \ldots, x_n$ be an ε-quasiapproximation to L. First we show, for each z in M, that $r = d(z, \{x_1, \ldots, x_n\}) - \varepsilon$ is a lower bound for $T = \{d(z, y) : y \in L\}$. Given y in L there is i such that $d(x_i, y) < \varepsilon$, so $d(z, x_i) \leq d(z, y) + d(y, x_i) < d(z, y) + \varepsilon$. Thus

$$r + \varepsilon \leq d(z, x_i) < d(z, y) + \varepsilon$$

so $r < d(z, y)$. Next we show that $r + 3\varepsilon$ is not a lower bound for T. We know that $d(z, x_i) < r + 2\varepsilon$ for some i, so $B_\varepsilon(x_i) \subset B_{r+3\varepsilon}(z)$. We also know that $B_\varepsilon(x_i) \cap L$

cannot be empty, so $B_{r+3\varepsilon}(z) \cap L$ cannot be empty, whence $r + 3\varepsilon$ is not a lower bound for T. Thus we can approximate $\delta(z, L)$ within $3\varepsilon/2$. □

We say that L is a **quasi-n-set** if for each $\varepsilon > 0$ there is an ε-quasiapproximation to L of length n. Note that a singleton is a quasi-n-set for every positive n, so n is just an upper bound on the size of L rather than its size. We say that L is **quasifinite** if it is a quasi-n-set for some n. Clearly a quasifinite set is quasi-totally-bounded.

The meaning of $\delta(x, L) < t$, when $\delta(x, L)$ does not necessarily exist, follows from the general theory of lowercuts. It is the second condition in the following lemma.

Lemma 2.2. *If $\delta(x, L)$ exists, then $\delta(x, L) < t$ if and only if $B_r(x) \cap L$ cannot be empty for some $r < t$.*

Proof. We know that $r \leq \delta(x, L)$ if and only if $B_r(x) \cap L$ is empty (Lemma 1.4), so if $\delta(x, L) < t$, then $B_r(x) \cap L$ cannot be empty for $r = (\delta(x, L) + t)/2$. Conversely, if $B_r(x) \cap L$ cannot be empty for some $r < t$, then $r \leq \delta(x, L)$ is impossible, so $\delta(x, L) < t$. □

We will show, in Corollary 4.6, that the zero set of a monic polynomial is quasifinite.

3. The spectrum of a polynomial

Let f be a monic polynomial of degree $n \geq 1$ with complex coefficients. We can approximate f arbitrarily closely with polynomials of the form

$$(X - q_1)(X - q_2) \cdots (X - q_n)$$

where the q_i are Gaussian numbers. The **spectrum** of f, written $\operatorname{spec} f$, is an element of the completion of the set of n-multisets of Gaussian numbers in an appropriate metric, see [3]. Classically, $\operatorname{spec} f$ can be identified with the multiset of roots of f. Constructively, with choice, $\operatorname{spec} f$ can be identified with an equivalence class of multisets $r_1, \ldots, r_n$ where two multisets are equivalent if $\prod_{i=1}^{n}(X - r_i) = \prod_{i=1}^{n}(X - r'_i)$.

The choiceless fundamental theorem of algebra says that the function spec is a homeomorphism from the space of monic polynomials of degree n, with the natural metric, to the completion of the set of n-multisets of Gaussian numbers.

The **distance**, $d(z, \operatorname{spec} f)$, from a complex number z to $\operatorname{spec} f$ is approximated by the distance from z to the approximating multisets $\{q_1, q_2, \ldots, q_n\}$. Note that $d(z, \operatorname{spec} f)$ is independent of multiplicities. It is not hard to show that $d(z, \operatorname{spec} f) = 0$ if and only if $f(z) = 0$. That is the connection between $\operatorname{spec} f$ and the zero set of f.

The **diameter** of $\operatorname{spec} f$ is the limit of the the diameters of its approximating n-multisets. Caution: this need not be a limit of a *sequence* of approximating n-multisets—we are operating in the absence of countable choice. The diameter can be characterized in terms of $d(z, \operatorname{spec} f)$ as

$$\inf_{\varepsilon > 0} \sup \{|z - z'| : d(z, \operatorname{spec} f) < \varepsilon \text{ and } d(z', \operatorname{spec} f) < \varepsilon\}$$

The (set-)**distance** between $\operatorname{spec} f$ and $\operatorname{spec} g$ is defined to be the infimum over all z of

$$d(z, \operatorname{spec} f) + d(z, \operatorname{spec} g).$$

It is not hard to see that this infimum exists, and is equal to the limit of the distance between the approximations to $\operatorname{spec} f$ and $\operatorname{spec} g$. This is like the distance between two sets of complex numbers. It is not a metric and should not be confused with the metric distance between the spectra of two monic polynomials of the same degree that comes from their both being in the completion of the set of n-multisets of Gaussian numbers.

We will need this lemma.

Lemma 3.1. *Let f be a monic polynomial of degree $n \geq 1$ with complex coefficients. If the diameter of* $\operatorname{spec} f$ *is positive, then $f = gh$ for nonconstant polynomials g and h. Moreover, the distance between* $\operatorname{spec} g$ *and* $\operatorname{spec} h$ *is positive.*

Proof. Let d be the diameter of $\operatorname{spec} f$ and let r be a rational number such that $d \geq r > 0$. Let $q_1, \ldots, q_n$ be Gaussian numbers that approximate the spectrum of f to within $r/(8n)$. The transitive closure of the relation $|x - y| \leq r/(2n)$ partitions the q_i into at least two equivalence classes (thought of as multisets). Let G_0 be one of them and H_0 the union of all the rest. The minimum distance ε between a point in G_0 and a point in H_0 is greater than $r/(2n)$.

Let m be the number of elements of G_0, including multiplicities. Any multiset of Gaussian numbers that approximates the spectrum of f to within $r/(8n)$ is uniquely the union of an m-multiset G and an $(m-n)$-multiset H, where G is within $r/(4n)$ of G_0, as m-multisets, and H is within $r/(4n)$of H_0. So the distance between any point in G and any point in H is at least $\varepsilon - r/(2n)$. Let g and h be the monic polynomials whose spectra are the limits of the multisets G and H respectively. That is, g is the limit of the polynomials $\Pi_{s \in G}(X - s)$. The distance between $\operatorname{spec} g$ and $\operatorname{spec} h$ is at least $\varepsilon - r/(2n)$. □

Any easy induction gives us the theorem.

Theorem 3.2. *Let f be a monic polynomial of degree $n \geq 1$ with complex coefficients. For each $\varepsilon > 0$ we can write $f = g_1 g_2 \cdots g_k$ where the diameter of each* $\operatorname{spec} g_i$ *is less than ε and the* $\operatorname{spec} g_i$ *are a positive distance from each other.*

Proof. If the diameter d of $\operatorname{spec} f$ is less than ε, then let $k = 1$ and $g_1 = f$. Otherwise $d > 0$ and from Lemma 3.1 we can write $f = gh$ for nonconstant polynomials g and h. By induction on degree we can write g and h in the desired form. Clearly the spectrum of any factor of g is a positive distance from the spectrum of any factor of h. □

4. Weak fundamental theorem of algebra

We want to show that the zero set L of a monic polynomial p is quasilocated. We know that the quasidistance from z to L has to be $r = d(z, \operatorname{spec} p)$. But to prove that, we must show that it is impossible for $B_{r+\varepsilon}(z) \cap L$ to be empty for $\varepsilon > 0$ (Lemma 1.4). In particular, we must show that it is impossible for the zero set of a monic polynomial of degree at least 1 over the complex numbers to be empty—the weak fundemental theorem of algebra.

Suppose we have two polynomials

$$\begin{aligned} a &= a_m X^m + a_{m-1} X^{m-1} + \cdots + a_1 X + a_0 \\ b &= b_n X^n + b_{n-1} X^{n-1} + \cdots + b_1 X + b_0 \end{aligned}$$

with coefficients in a commutative ring R. The so-called **pseudodivision algorithm**, which is proved in essentially the same way as the division algorithm for monic polynomials, says that if $0 < m \leq n$, then we can find canonical q and r, with $\deg r < m$, such that

$$a_m^{n-m+1} b = qa + r$$

We will denote the remainder r by $\rho(b, a)$ and we say that $\rho(b, a)$ is the result of **remaindering** b by a.

The coefficients of q and r can be written as polynomials in the coefficients of a and b. The computation really takes place in the polynomial ring

$$\mathbf{Z}[\alpha_m, \ldots, \alpha_0, \beta_n, \ldots, \beta_0][X]$$

where the α's and β's are indeterminates, so we might as well assume that R is discrete. Note that $\rho(b, a)$ is defined only when $0 \leq \deg a < \deg b$.

Let R be a discrete commutative ring. If a is any nonzero polynomial in $R[X]$, then $\gamma(a)$ is defined to be the result of deleting the leading term of a. We say that $\gamma(a)$ is the **chop** of a. Set $\gamma(0) = 0$.

If A and B are subsets of $R[X]$, then $\operatorname{rem}(A, B)$ is defined to be

$$\{p \in R[X] : p = \rho(a, b) \text{ or } p = \rho(b, a) \text{ for some } (a, b) \in A \times B\}$$

and $\operatorname{chop}(A)$ is defined to be $\{\gamma(a) : a \in A\}$. We say that A is **closed under chopping and remaindering** if $\operatorname{rem}(A, A) \subset A$ and $\operatorname{chop}(A) \subset A$.

Lemma 4.1. *Given a finite set S of polynomials over a discrete commutative ring, there is a finite set of polynomials containing S that is closed under chopping and remaindering.*

Proof. Define S_0 to be S. Let n be a bound on the degrees of the polynomials in S_0. For $i > 0$, define

$$S_i = S_{i-1} \cup \operatorname{chop}(S_{i-1}) \cup \operatorname{rem}(S_{i-1}, S_{i-1})$$

We will show that, for $i \geq 1$, the elements of $S_i \setminus S_{i-1}$ have degree at most $n - i$. That's clearly true for $i = 1$, and if $i > 1$, and $p \in S_i \setminus S_{i-1}$, then

$$p \in \operatorname{chop}(S_{i-1} \setminus S_{i-2}) \cup \operatorname{rem}(S_{i-1}, S_{i-1} \setminus S_{i-2})$$

whence, by induction on i, we have $\deg p < n - (i-1)$ so $\deg p \leq n - i$. In particular, the elements of $S_n \setminus S_{n-1}$ have degree at most 0, so S_n is closed under chopping and remaindering. □

Apply Lemma 4.1 to the set S of two polynomials

$$\begin{aligned} &\alpha_m X^m + \alpha_{m-1} X^{m-1} + \cdots + \alpha_1 X + \alpha_0 \\ &\beta_n X^n + \beta_{n-1} X^{n-1} + \cdots + \beta_1 X + \beta_0 \end{aligned}$$

over the ring $\mathbf{Z}[\alpha_m, \ldots, \alpha_0, \beta_n, \ldots, \beta_0]$, where the α's and β's are indeterminates. Let L be the set of leading coefficients of polynomials in S'. The elements of L are

in $\mathbf{Z}[\alpha_m, \ldots, \alpha_0, \beta_n, \ldots, \beta_0]$. If we are then given two polynomials

$$\begin{aligned} a &= a_m X^m + a_{m-1} X^{m-1} + \cdots + a_1 X + a_0 \\ b &= b_n X^n + b_{n-1} X^{n-1} + \cdots + b_1 X + b_0 \end{aligned}$$

with coefficients in an arbitrary commutative ring R, we can form the subset $L(a, b)$ of R obtained by plugging a_i and b_j for α_i and β_j in the elements of L. We are now ready to prove a general Bezout's equation for polynomials over an arbitrary commutative ring.

Theorem 4.2. *Let*

$$\begin{aligned} a &= a_m X^m + a_{m-1} X^{m-1} + \cdots + a_1 X + a_0 \\ b &= b_n X^n + b_{n-1} X^{n-1} + \cdots + b_1 X + b_0 \end{aligned}$$

be polynomials over a commutative ring R. If each of the elements of $L(a, b)$ is either 0 or invertible, then there exist s and t in $R[X]$ such that $sa + tb$ divides both a and b and is either monic or 0.

Proof. We attempt to execute the Euclidean algorithm on a and b. This will succeed if, each time we create a new remainder, we can write it as a polynomial with an invertible leading coefficient. But that is exactly what the hypothesis of the theorem allows us to do. The elements of $L(a, b)$ include all the coefficients of a and b and anything we might get by repeatedly applying the division algorithm. Unless $a = b = 0$, the last nonzero remainder from the algorithm is of the form $sa + tb$ and divides both a and b. We can divide by its leading coefficient to make it monic. □

As an immediate corollary we get

Corollary 4.3. Let R be a ring in which any noninvertible element is zero. If a and b are two polynomials in $R[X]$ then we can construct propositions $P_1, \ldots, P_k$ such that if the conjunction

$$(P_1 \vee \neg P_1) \wedge (P_2 \vee \neg P_2) \wedge \cdots \wedge (P_k \vee \neg P_k)$$

of instances of the law of excluded middle holds, then there exist s and t in $R[X]$ such that $sa + tb$ divides both a and b and is either monic or 0.

The complex numbers are such a ring, as is any Heyting field. Why do we want Corollary 4.3? The relevant intuitionistic tautology is that the negation of

$$(P_1 \vee \neg P_1) \wedge (P_2 \vee \neg P_2) \wedge \cdots \wedge (P_k \vee \neg P_k) \wedge A$$

implies $\neg A$. That is easy to see by induction from the case $k = 1$. For $k = 1$, suppose A and $\neg((P_1 \vee \neg P_1) \wedge A)$. If P_1 is true, then we get a contradiction, so P_1 is false, and we get a contradiction. So if we want to prove $\neg A$, we are allowed to assume any finite number of instances of the law of excluded middle. The proposition A that we are interested in refuting is that a given monic polynomial with complex coefficients has no root. That refutation is the weak fundamental theorem of algebra.

First we prove the fundamental theorem of algebra for strongly separable polynomials.

Theorem 4.4. *Let f be a monic polynomial with complex coefficients. Suppose that there exist polynomials s and t such that $sf+tf'=1$. Then $f=(X-r_1)\cdots(X-r_n)$ where the r_i are distinct complex numbers.*

Proof. From [3] we can find algebraic numbers $q_1,\ldots,q_n$ that approximate the spectrum of f arbitrarily closely. That means that

$$g=(X-q_1)\cdots(X-q_n)$$

is arbitrarily close to f, and that given another close approximation $q_1',\ldots,q_n'$ to the spectrum of f, we can reindex it so that q_i' is close to q_i for each i. In particular, the numbers $f(q_i)$ can be made arbitrarily small so the numbers $f'(q_i)$ are bounded away from 0 because of the equation $sf+tf'=1$. Differentiation of polynomials is continuous, so $f'(q_1)$ is close to

$$g'(q_1)=(q_1-q_2)(q_1-q_3)\cdots(q_1-q_n).$$

Hence q_1 is bounded away from $q_2,\ldots,q_n$, and similarly for the rest of the q's. Thus we can find n discs, $D_1,\ldots,D_n$, bounded away from each other, so that any sufficiently close approximation to the spectrum of f has one point in each of the discs. The points in D_i, as we range over approximations to the spectrum, form a coherent system of approximations, hence determine a unique complex number r_i because the complex numbers are complete in this choiceless sense. Clearly f vanishes on each r_i, so f can be written as desired. □

Now we can prove the weak fundamental theorem of algebra for arbitrary monic polynomials.

Theorem 4.5 (Weak fundamental theorem of algebra)**.** *Let f be a nonconstant monic polynomial with complex coefficients. Then the assumption that f has no roots leads to a contradiction.*

Proof. We apply Corollary 4.3 to the polynomial f and its derivative f'. Assume f has no roots. We will derive a contradiction from that and the existence of polynomials s and t in $\mathbf{C}[X]$ such that $d=sf+tf'$ is a monic polynomial that divides both f and f'. As f' is a nonzero polynomial of degree less than f, the polynomial d is either 1 or a monic nonconstant proper divisor of f. In the latter case, we can replace f by d and we are done by induction on the degree as d cannot have roots (if $\deg f=1$, then f has a root). So we may assume that

$$sf+tf'=1.$$

But then Theorem 4.4 says that f has a root. □

The following corollary is the point of this whole exercise.

Corollary 4.6. Let f be a nonconstant monic polynomial with complex coefficients, Z the zero set of f, and z a complex number. If $d(z,\operatorname{spec} f)=r$, and $\varepsilon>0$, then the assumption that f has no roots in the disc of radius $r+\varepsilon$ around z leads to a contradiction. Thus $d(z,\operatorname{spec} f)=\delta(z,Z)$ so Z is quasilocated. In fact, Z is quasifinite.

Proof. Write $f=g_1g_2\cdots g_k$ where the diameter of each $\operatorname{spec} g_i$ is less than ε. As $r=\inf_i d(z,\operatorname{spec} g_i)$, there is i such that the diameter of g_i plus $d(z,\operatorname{spec} g_i)$ is

less than $r+\varepsilon$. That requires that g_i has no roots (anywhere), which leads to a contradiction by Theorem 4.5.

To show that Z is quasifinite, we have to find, for each $\varepsilon > 0$, an ε-quasi-approximation to Z of length n. An ε-quasi-approximation to Z is given by the root set $r_1, \ldots, r_n$ of a polynomial over the Gaussian numbers that is close to f. To show that it is an ε-quasi-approximation, we need to show that $B_\varepsilon(r_i) \cap Z$ cannot be empty for $i = 1, \ldots, n$, and $Z \subset B_\varepsilon(r_1) \cup \cdots \cup B_\varepsilon(r_n)$. The first condition means that $\delta(r_i, Z) < \varepsilon$. But $\delta(r_i, Z) = d(r_i, \operatorname{spec} f)$ which we can make arbitrarily small. For the second condition, if $z \in Z$, so $f(z) = 0$, then z is close to one of the r_i because if the polynomial f is close to $(X - r_1) \cdots (X - r_n)$, and $f(z) = 0$, then z is close to one of the r_i. □

5. Comaximal polynomials and resultants

We would like to prove that, for monic polynomials a and b over $\mathbf{C}$, that the distance between $\operatorname{spec} a$ and $\operatorname{spec} b$ is positive if and only if there are polynomials s and t such that $sa + tb = 1$. To do this, we will show that if the distance from $\operatorname{spec} a$ to $\operatorname{spec} b$ is positive, then the resultant of a and b is different from zero. Then we will apply a theorem valid for arbitrary commutative rings (Theorem 5.1) that will construct the polynomials s and t.

The **resultant** of a and b (not necessarily monic) is the determinant of their Sylvester matrix. Here is the **Sylvester matrix** for $a = a_3X^3 + a_2X^2 + a_1X + a_0$ and $b = b_5X^5 + \cdots + b_1X + b_0$

$$S = \begin{pmatrix} a_3 & a_2 & a_1 & a_0 & 0 & 0 & 0 & 0 \\ 0 & a_3 & a_2 & a_1 & a_0 & 0 & 0 & 0 \\ 0 & 0 & a_3 & a_2 & a_1 & a_0 & 0 & 0 \\ 0 & 0 & 0 & a_3 & a_2 & a_1 & a_0 & 0 \\ 0 & 0 & 0 & 0 & a_3 & a_2 & a_1 & a_0 \\ b_5 & b_4 & b_3 & b_2 & b_1 & b_0 & 0 & 0 \\ 0 & b_5 & b_4 & b_3 & b_2 & b_1 & b_0 & 0 \\ 0 & 0 & b_5 & b_4 & b_3 & b_2 & b_1 & b_0 \end{pmatrix}$$

If $a = \prod_{i=1}^{m}(X - q_i)$ and $b = \prod_{j=1}^{n}(X - r_j)$, then the determinant of S is $\prod_{i,j}(q_i - r_j)$. That fact is a simply a polynomial identity in the indeterminants q_i and r_j.

We say that two elements a and b a ring are **comaximal** if there exist s and t in the ring such that $sa + tb = 1$.

Theorem 5.1. *Let $a = a_mX^m + a_{m-1}X^{m-1} + \cdots + a_0$ and $b = b_nX^n + b_{n-1}X^{n-1} + \cdots + a_0$ be polynomials in $R[X]$ with $m, n \geq 1$. Then the following two conditions are equivalent:*

1. *The polynomials a and b are comaximal in $R[X]$ and the (formal) leading coefficients a_m and b_n are comaximal in R,*
2. *The resultant of a and b is a unit in R.*

Proof. If S is the Sylvester matrix of a and b, and S^* is the adjugate of S, then $S^*S = (\det S)\, I$, where I is the identity matrix and $\det S$ is the resultant of a and b. So the last row of S^* gives the coefficients of polynomials s and t such that $sa + tb$ is the resultant of a and b. Thus if condition 2 holds, then the polynomials a and b are comaximal in $R[X]$. It remains to show that a_m and b_n are comaximal in

R. But that follows by expanding the determinant of S by minors using the first column.

Now suppose condition 1 holds, so there exist polynomials s and t such that $sa + tb = 1$. We will first prove 2 under the assumption that a_m (or b_n) is a unit. Now if g is any polynomial of degree less than $m + n$, then we can certainly find polynomials s and t such that $sa + tb = g$. As a_m is unit, we can write $t = qa + r$ where $\deg r < m$. So

$$(s + qb)\, a + rb = g$$

which means we may assume that $\deg t < m$, so $\deg tb < m + n$. This forces $\deg s < n$ because if s_k is the (formal) leading coefficient of s, and $k \geq n$, then the equation $sa + tb = g$ requires $s_k = 0$ as $\deg tb < m + n$. Because $\deg s < n$ and $\deg t < m$, we can put the coefficients of s and t together to form a row vector v of length $m + n$ so that $vS = (g_{m+n-1}, g_{m+n-2}, \ldots, g_1, g_0)$ where S is the Sylvester matrix. Since we can do this for arbitrary g of degree less than $m + n$, this shows that the matrix S is invertible, whence $\det S$, which is the resultant of a and b, is a unit in R.

To finish the proof, we pass to the rings $R[1/a_m]$ and $R[1/b_n]$. The kernel of the natural map $R \to R[1/a_m]$ is the set of elements of R that annihilate some power of a_m. In particular, $R[1/a_m]$ is trivial exactly when a_m is nilpotent. Now since the theorem is true when either a_m or b_n is a unit, it follows that $\det S$ is a unit in $R[1/a_m]$ and in $R[1/b_n]$. We want to show that $\det S$ is a unit in R. That follows from the fact that a_m and b_n are comaximal. Specifically, we have $(\det S)\, r/a_m^k = 1$ in $R[1/a_m]$ for some $r \in R$ and positive integer k, so $(\det S)\, r a_m^i = a_m^{k+i}$ in R for some positive integer i. Thus $\det S$ divides a power of a_m in R. Similarly, $\det S$ divides a power of b_n. That makes $\det S$ a unit because any power of a_m is comaximal with any power of b_n. □

It's not enough in Theorem 5.1 just to require that a and b be comaximal. A trivial example is $a = b = 0X + 1$ where $m = n = 1$.

Over a discrete field, Theorem 5.1 is often stated as the resultant of a and b is zero if and only if a and b have a nontrivial common factor, it being assumed that nobody is so silly as to have the formal leading coefficient of a and b equal to zero. The theorem in this form, even for monic a and b, does not generalize well to arbitrary commutative rings.

Consider the polynomials $(X - 2)(X - 6) = X^2 + 4$ and $X(X - 4)$ over $\mathbf{Z}_8$. It is easily checked that their resultant is zero. However, they have no nontrivial common factors. It's easy to see that they have no nontrivial *monic* common factors. What constitutes a trivial common factor? It seems clear that these are exactly the unit polynomials, the polynomials with inverses. Over a decent ring, the units in $R[X]$ are just the units of R, but in a ring with nilpotent elements, like $\mathbf{Z}_8$, there are units of arbitrarily high degree, like $4X^n + 1$.

We will show that over $\mathbf{Z}_8$, or any primary ring, the factors of a monic polynomial are all products of a unit polynomial and a monic polynomial. Thus if a monic polynomial has a nontrivial factor, then it has a nontrivial monic factor.

First we note that in any ring, the factorization of a polynomial as a monic polynomial times a unit polynomial, if possible, is unique. Indeed, if $mu = m'u'$, then $m = m'u'u^{-1}$ so $\deg m' \leq \deg m$, hence, by symmetry $\deg m' = \deg m$. As both m and m' monic, we must have $u'u^{-1} = 1$.

We can characterize the polynomials that can be factored in this way.

Theorem 5.2. *Let $p = \sum a_i X^i$ be a polynomial over a commutative ring R. Suppose a_m is a unit in R and that a_i is nilpotent for $i > m$. Then p can be written as a unit polynomial times a monic polynomial.*

Proof. Let I be a fixed nilpotent ideal of R such that $a_i \in I$ for all $i > m$. We will multiply p by a sequence of unit polynomials until it becomes monic.

Suppose $j > m$ and $d \geq 1$ are such that

$$a_i \in I^d \text{ for } m < i \leq j$$
$$a_i \in I^{d+1} \text{ for } i > j$$

Multiply p by the unit polynomial

$$1 - \frac{a_j}{a_m} X^{j-m}$$

In the resulting polynomial p' we have

$$a'_m \in a_m + I^d \text{ is a unit}$$
$$a'_j = 0 \in I^{d+1}$$
$$a'_i \in I^d + I^d = I^d \text{ for } m < i \leq j$$
$$a'_i \in I^{d+1} + I^{d+d} = I^{d+1} \text{ for } i > j$$

When $j = m + 1$, we get $a'_i \in I^{d+1}$ for all $i > m$. At that point we set d equal $d + 1$ and continue until $I^d = 0$. At that point the polynomial has a unit a_m as its formal leading coefficient. □

The converse of Theorem 5.2 is also true because the units in the ring $R[X]$ are exactly those polynomials whose constant term is a unit in R and whose other coefficients are nilpotent.

We end this section by applying Theorem 5.1 to polynomials over the complex numbers and their spectra.

Theorem 5.3. *If a and b are nonconstant monic polynomials over $\mathbf{C}$, then the distance between* spec a *and* spec b *is positive if and only if a and b are comaximal.*

Proof. We will apply Theorem 5.1. As a and b are monic, their leading coefficients are certainly comaximal. So a and b are comaximal if and only if their resultant is a unit, that is, if and only if the resultant is different from zero. It thus suffices to show that the resultant is different from zero if and only if the distance between spec a and spec b is positive.

The resultant, as a function of the coefficients of a and b, is clearly uniformly continuous on bounded subsets. So if $\prod_{i=1}^{m} (X - q_i)$ and $\prod_{j=1}^{n} (X - r_j)$ are close to a and b respectively, then $\prod_{i,j} (q_i - r_j)$ is close to the resultant of a and b. Thus the resultant of a and b is the limit of $\prod_{i,j} (q_i - r_j)$. On the other hand, the distance between spec a and spec b is the limit of $\min |q_i - r_j|$. Clearly $\prod_{i,j} (q_i - r_j)$ is eventually bounded away from zero if and only if $\min |q_i - r_j|$ is. □

6. Examples of Riesz spaces

The question is raised in [2] as to whether you can construct the points of the spectrum of a separable Riesz space, all of whose elements are normable, in the absence of countable choice. The authors, referring to [3], suggested basing a counterexample on the zero set of $X^2 - a$. However it appeared that the resulting Riesz space was not separable. Here we elaborate on that construction and give a separable counterexample.

A **Riesz space** is a lattice ordered vector space V over the rational numbers with an element 1 such that for each v in V, there exists a positive integer n such that $v \leq n \cdot 1$. The canonical example is the space of uniformly continuous functions on a compact metric space, where 1 is the constant function whose value is 1. We say that V is **normable**, or **normed**, if for each $v \in V$, the set $\{q \in \mathbf{Q} : v \leq q \cdot 1\}$ has an infimum in $\mathbf{R}$. We call this greatest lower bound the **least upper bound** of v. The norm of f is defined to be the least upper bound of $|v| = v \vee 0 - v \wedge 0$. Note that classically every Riesz space is normed.

Let S be a quasilocated subset of the complex plane that is contained in a closed disc D. Let $C(D)$ be the Riesz space of uniformly continuous real valued functions on D. Let K be the functions in $C(D)$ that vanish on S. The set K is a subspace of $C(D)$ that is closed under the lattice operations. Our counterexample is the space $C(D)/K$. It is separable because $C(D)$ is. Note that this is not exactly the space of uniformly continuous functions on S, it is the space of those uniformly continuous functions on S that extend to uniformly continuous functions on D.

If $f \in C(D)$, then f vanishes on S if and only if f and $-f$ are nonnegative on S. Note that if f is nonnegative on S, and g vanishes on S, then $f + g$ is nonnegative on S, so this defines a notion of nonnegativity on $C(D)/K$. You get another notion of nonnegativity on $C(D)/K$ by considering those functions that come from nonnegative functions in $C(D)$. However, if f is nonnegative on S, then f is equal to $f \vee 0$ on S, so these two notions are the same.

Theorem 6.1. *Let S be a quasilocated subset of the complex plane that is contained in a closed disc D. Let $C(D)$ be the Riesz space of uniformly continuous real valued functions on D and let K be the functions in $C(D)$ that vanish on S. Then every element of $C(D)/K$ is normable, and each Riesz-space homomorphism $C(D)/K \to R$ that takes 1 to 1 is given by evaluation at a point $z_0 \in D$ such that $\delta(z_0, S) = 0$. Conversely, if $\delta(z_0, S) = 0$, and $f \in K$, then $f(z_0) = 0$, so evaluation at z_0 gives a Riesz-space homomorphism $C(D)/K \to R$ that takes 1 to 1.*

Proof. We want to approximate the least upper bound of f on S to within ε. Choose $\theta > 0$ so that if $d(x, y) \leq \theta$, then $|f(x) - f(y)| < \varepsilon/2$. Take a finite $\theta/4$-approximation F to D and partition F into two subsets, one, S', where $\delta(x, S) < \theta$ and one where $\delta(x, S) > \theta/2$. Let μ be the supremum of f on S'. We want to show that $f \leq \mu + \varepsilon$ on S and that it is not the case that $f \leq \mu - \varepsilon$ on S.

To show that $f \leq \mu + \varepsilon$ on S, let $s \in S$. There is $s' \in S'$ such that $d(s, s') < \theta$ for otherwise $d(s, s') > \theta/2$ for every $s' \in S'$. But $d(s, t) < \theta/2$ for some $t \in F$, and this t must be in S'. Thus $|f(s) - f(s')| < \varepsilon/2$. As $f(s') \leq \mu$, we conclude that $f(s) \leq \mu + \varepsilon/2$.

Now suppose $f \leq \mu - \varepsilon$ on S. We know that there is $s' \in S'$ such that $f(s')$ is arbitrarily close to μ. Note that $B_\theta(s') \cap S$ cannot be empty because $\delta(s', S) < \theta$. Yet if $s \in B_\theta(s') \cap S$, then $|f(s) - f(s')| < \varepsilon/2$, so $f(s') \leq \mu - \varepsilon/2$, a contradiction.

For the second claim, let ρ be a Riesz-space homomorphism $C(D)/K \to \mathbf{R}$ taking 1 to 1. This induces a homomorphism $\rho : C(D) \to \mathbf{R}$ taking 1 to 1 and K to zero. To see that ρ is given by evaluation at a point in D, let W be the vector sublattice of $C(D)$ generated by 1 and the two coordinate projections, π_1 and π_2. Then W is dense in $C(D)$ by the Stone-Weierstrass approximation theorem (see below). Let $z_0 = \rho(\pi_1) + i\rho(\pi_2)$. It's easy to see that ρ is evaluation at z_0 on W, hence is evaluation at z_0 on $C(D)$.

We want to show that $\delta(z_0, S) = 0$. The function $f(x) = \delta(x, S)$ is in $C(D)$ and vanishes on S. So $\rho(f) = 0$. But ρ is evaluation at z_0, so $\delta(z_0, S) = 0$. Conversely, suppose $\delta(z_0, S) = 0$ and $f \in K$. Then Lemma 1.7 tells us that $f(z_0) = 0$. □

Corollary 6.2. Let S be the root set of a nonconstant monic polynomial p and let $V = C(D)/K$ be the Riesz space associated with S as in the theorem. Then V is normed and separable, and each Riesz-space homomorphism from V onto $\mathbf{R}$ gives a root of p.

Proof. The space V is separable because the space $C(D)$ is separable. As p is in K, if $\delta(z_0, S) = 0$, then z_0 is a root of p. □

The point of the corollary is that, in the absence of choice, we need not be able to construct a root of the polynomial $p(X) = X^2 - a$, see [3], so we need not be able to construct a Riesz-space homomorphism from a normed, separable, Riesz space onto $\mathbf{R}$.

We include a short choice-free proof of the Stone-Weierstrass theorem for the square, which was appealed to in the proof of Theorem 6.1. We want to approximate nonegative functions on the square by finite suprema of finite infima of functions of the form $ax + by + c$, that is, functions whose graphs are planes. Suppose we have a little square S_i in a grid and a nonnegative constant c_i. Let φ_i be the infimum of the constant c_i and one affine function for each side L of S_i whose value is c_i on L and 0 on the opposite side of the other little square in the grid with side L. The graph of φ_i is a pyramid with its top cut off. Let $\psi = \sup_i \varphi_i$. Let J be the indexes of the block of nine little squares centered at S_i. If $x \in S_i$, then

$$\inf_{j \in J} c_j \leq \psi(x) \leq \sup_{j \in J} c_j$$

Given a uniformly continuous function f on the square, choose a grid so that f does not vary by more than ε on any block of nine little squares in the grid. Choose c_i to be the value of f in the center of the little square S_i and let ψ be as above. Then ψ approximates f within ε.

References

[1] Bishop, Errett, Foundations of constructive analysis, McGraw-Hill 1967

[2] Coquand, Thierry and Bas Spitters, Formal topology and constructive mathematics: the Gelfand and Stone-Yoshida representation theorems, *Journal of Universal Computer Science*, **11** (2005) 1932–1944

[3] Richman, Fred, The fundamental theorem of algebra: a constructive development without choice, *Pacific Journal of Mathematics*, **196** (2000), 213-230. MR 2001k:03141

[4] _____, Generalized real numbers in constructive mathematics, *Indagationes Mathematicae*, **9** (1998), 595-606. MR 2000e:03172

[5] Ruitenburg, Wim B. G., Constructing roots of polynomials over the complex numbers. *Computational aspects of Lie group representations and related topics* (Amsterdam, 1990), 107–128, CWI Tract, **84**, Math. Centrum, Centrum Wisk. Inform., Amsterdam, 1991.

Florida Atlantic University
E-mail address: Robert.Lubarsky@comcast.net

Florida Atlantic University
E-mail address: richman@fau.edu